



The
Patent
Office

09/806968 PCT/GB 99 / 03356



GB99/3356

ESU

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.


**PRIORITY
DOCUMENT**

SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH RULE 17.1(a) OR (b)

Signed

Dated


19 NOV 1999



THIS PAGE BLANK (uspto)

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)



The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference

PDG/20952

2. Patent application number

(The Patent Office will fill in this part)

09 OCT 1998

9822094.0

3. Full name, address and postcode of the or of each applicant (underline all surnames)

SNELL & WILCOX LIMITED

6 Old Lodge Place

St Margaret's

Twickenham

Middlesex TW1 1RQ

557978403

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

UNITED KINGDOM

4. Title of the invention

IMPROVEMENTS IN DATA COMPRESSION

5. Name of your agent (if you have one)

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

MATHYS & SQUIRE

100 Gray's Inn Road

London WC1X 8AL

UNITED KINGDOM

Patents ADP number (if you know it)

1081001

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

Yes

- a) any applicant named in part 3 is not an inventor, or
 - b) there is an inventor who is not named as an applicant, or
 - c) any named applicant is a corporate body.
- See note (d))

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form	-
Description	10
Claim(s)	-
Abstract	-
Drawing(s)	-

10. If you are also filing any of the following, state how many against each item.

Priority documents	-
Translations of priority documents	-
Statement of inventorship and right to grant of a patent (<i>Patents Form 7/77</i>)	-
Request for preliminary examination and search (<i>Patents Form 9/77</i>)	-
Request for substantive examination (<i>Patents Form 10/77</i>)	-
Any other documents (<i>please specify</i>)	-

11.

I/We request the grant of a patent on the basis of this application.

Signature

MATHYS & SQUIRE

9 Sep 1998

Date

12. Name and daytime telephone number of person to contact in the United Kingdom

Peter D Garratt - 0171 830 0000

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

IMPROVEMENTS IN DATA COMPRESSION

1 Introduction

This patent proposal describes methods for automatically estimating the subjective quality of a picture that has been decoded from a compressed bitstream. It is assumed that both the bitstream itself and the decoded picture are accessible but that the original source is not available, hence the term 'single-ended'. Such an estimate will clearly not be as reliable as one in which the source picture can be compared to the decoded output, but it can serve as a useful indicator of potential problems in a broadcast chain involving compression when the bitstream is being monitored.

2 Terminology and prior art

This patent relates to a hybrid transform based video compression system. The classic example is the MPEG-2 video compression standard [1].

The problem to be solved is that of estimating the subjective picture quality of a picture or sequence decoded from an MPEG-2 bitstream. The usual method of performing such an estimate is referred to in this proposal as the "double-ended" method, as illustrated here:

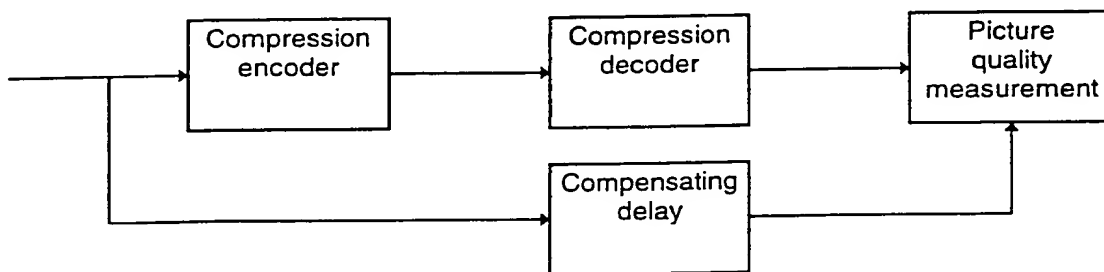


Figure 1 Double-ended quality measurement

The decoded picture is compared with a necessarily delayed version of the source picture. The most common quality measure based on this comparison is the peak signal-to-noise ratio (PSNR) which is based on the ratio of the maximum possible signal power to the power of the difference between source and decoded signals. Other measures are more sophisticated, for example, the one based on "Just Noticeable Differences" (JND) from Sarnoff Labs [2].

The disadvantage of all the methods based on the approach of Figure 1 is that they require access to the picture source. While this is appropriate for testing systems in a laboratory, it cannot normally be used for monitoring the quality of compression in the field. The object of the present invention is to overcome that disadvantage by providing a series of quality estimation methods based on a "single-ended" approach.

The single-ended approach makes use of the "Information Bus" which is the subject of an earlier patent application [3]. The Information Bus is a signal containing all the compression coding decisions and parameters extracted from the compressed bitstream, in an easily accessible form. More sophisticated versions of the quality estimation techniques presented here may also make use of the "Video Postmark" which is also the subject of an earlier patent application [4]. The Video Postmark is similar to the Information Bus but carries information about other processes that may have taken place upstream of the compression codec under consideration.

3 Description of the invention

3.1 Basic architecture

The basic architecture of single-ended quality measurement is shown here:




Figure 2 Architecture of single-ended quality estimation

The picture quality measurement process operates only from information available at the decoder side of the compression codec; the decoded video signal and the Information Bus containing the coding decisions and parameters. It has no access to the picture source. Because of this, the quality measurement can never be completely reliable because there is no way of telling which degradations in the picture are due to the current coding process and which were on the source. So it is not intended as a replacement for laboratory measurements based on the double-ended approach. But it is useful for monitoring applications in the field where a simple automatic indication of the "red - amber - green" variety is required. However, a modification by which some account can be taken of the source is described in Appendix A.

The remainder of this paper outlines specific measures based on MPEG-2 coding.

3.2 Blockiness

One of the most frequent complaints about MPEG-2 coded pictures is that they appear "blocky", meaning that the block and macroblock structure of the picture is visible. These blocking artefacts can occur for several reasons:

- Variation in quantizer scale between macroblocks
- Coarse quantization of DC coefficients in non-intra macroblocks
- Residual visibility of a prediction error resulting from a non-uniform motion vector field

Instead of attempting to analyse each of those possible causes, the "blockiness" measure proposed here is based simply on the end result, i.e. the decoded picture. There are various possible measures of blockiness, but the principle behind all of them is to compare pixel differences across block boundaries with pixel differences not across block boundaries. In the discussion that follows, care should be taken to recognize the distinction between *macroblock* (16x16 block) boundaries and *DCT block* (8x8 block) boundaries.

The following is an example of a measure of blockiness that works on macroblock boundaries:

- *Horizontal macroblockiness* = the picture-by-picture mean absolute horizontal adjacent luminance pixel difference across macroblock boundaries, expressed as a fractional increase over the mean absolute horizontal adjacent pixel difference not across DCT block boundaries

An example showing how this measure could be implemented in hardware is given in the following diagram:

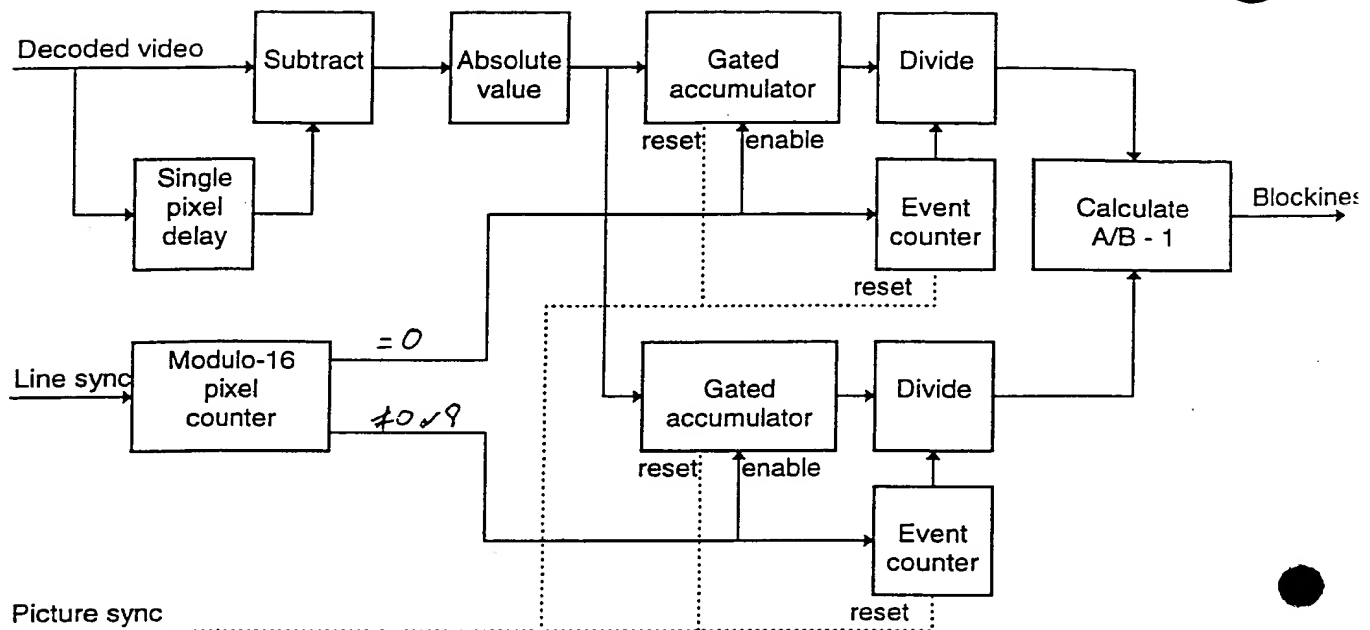


Figure 3 Architecture of horizontal macroblockiness measure

Pixel differences are taken across a pixel delay and the absolute value calculated. The result is fed to two gated accumulators controlled by a modulo-16 pixel counter which is reset by a line synchronization pulse. The upper accumulator sums the pixel differences across macroblock boundaries (when the modulo-16 pixel count = 0) and the lower accumulator sums the pixel differences not across DCT block boundaries (when the modulo-16 pixel count $\neq 0$ or 8). Event counters count the occurrences of each of these two cases so that the dividers can calculate mean values of the two quantities. Finally, the fractional increase is calculated, giving the blockiness measure. The accumulators and event counters are reset once per picture.

This particular measure has the interesting property that, when applied to frames that were I-frames in the MPEG-2 bitstream, the result is almost exactly proportional to the average quantizer scale value. When applied to P and B-frames, the result is smaller but reflects quite clearly differences in perceived blockiness arising from differences in motion estimation systems.

The following variations in the definition of the blockiness measure are possible and are considered to be part of the invention:

- The DCT block boundary can be used instead of the macroblock boundary. This would require a change the logical outputs of the pixel counter. Note that in both this and the original case the denominator of the fraction is the pixel difference not across DCT block boundaries.
- The difference could be taken vertically rather than horizontally (requiring a line delay instead of a pixel delay), or a combination of the two could be used. We have chosen the horizontal difference because this is much easier

to calculate in hardware and because the boundaries are the same whether field or frame picture coding was used. However, they may be circumstances in which the vertical differences are easier to calculate.

- Mean square values, or the mean of some other function of pixel differences, could be used instead of mean absolute differences.
- Some statistical function other than the mean could be used. For example, because it might be considered that very poor blockiness in a small region of the picture might be more disturbing to the eye than an evenly distributed blockiness resulting in the same average value, it might be better to use, for example, the 90th centile of the macroblock boundary pixel difference. A simple method for approximating such a statistical measure is given in Appendix B and should be considered part of this invention.
- The blockiness could be expressed as a logarithmic ratio (like a dB measure) rather than a fractional increase. This would affect the final block in Figure 3.
- It may be possible to use a reduced number of pixel differences in the measure.
- The measurement period could be greater or less than one picture period. This would affect the resetting of the accumulators and event counters in Figure 3.

In all cases it is necessary to record the blockiness separately for I-frames, P-frames and B-frames. The figures are much lower in P and B-frames because the denominator of the expression contains prediction residues that may have come from macroblock or block boundaries in reference frames. To detect the picture type (I, P or B), the Information Bus could be used. Alternatively, in the absence of the Information Bus, a method of picture type detection such as that described in [5] could be used. A further possibility is that the variations in the blockiness measure itself could be used as the basis of a method of picture type detection.

The above description assumes that the positions of the macroblock boundaries are known. In some cases, this information may not be available. However, it is possible to obtain this information by calculating the blockiness assuming each of the 16 possible positions in turn (either in full or using a reduced number of pixels) and choosing the position that yields the maximum value.

3.3 Quantizer consistency measure

A second single-ended quality measure is called the "quantizer consistency measure". This makes use only of the quantizer scale information and the picture type information, both of which are available on the Information Bus.

The principle behind the measure is to estimate variations in quantizing noise between different picture types, which might lead to visibility of the GOP structure.

The measure is defined by the following expression, which is calculated once per GOP:

$$\frac{1}{N} \sum N_k |q_k - \bar{q}|$$

where

$$\bar{q} = \frac{1}{N} \sum N_k q_k$$

N = number of frames in GOP

N_k = number of frames of type k in GOP

q_k = average q_scale_code in frames of type k

It is a weighted average of the variation in the average quantizer scale for each picture type from the overall average quantizer scale. A block diagram is not given because this measure would most likely be calculated in software.

3.4 Noise estimation based on quantizer scale

A third single-ended quality measure provides an estimate of the peak-signal-to-noise-ratio (PSNR) of the decoded picture using the quantizer scale values present in the bitstream.

The following is an outline of the theory by which quantizer scale values, in conjunction with some other parameters of the bitstream or of the decoded picture, may be used to estimate PSNR. The assumption is that the noise being estimated is the noise due to DCT coefficient quantization in the MPEG-2 coding process.

It is well known that the quantization noise power from a linear quantizer of spacing q in a uniformly distributed signal is given by the expression

$$\frac{q^2}{12}$$

In MPEG-2 coding, the signal consists of DCT coefficients and (apart from the intra DC coefficient) has a highly non-uniform probability distribution. A commonly accepted model of the probability distribution of DCT coefficients is the Laplacian distribution, which assuming a continuous variable, is given by the formula

$$p(x) = \frac{\alpha}{2} e^{-\alpha|x|}$$

If a signal has such a Laplacian distribution and is quantized using a quantizer with uniformly spaced levels apart from a decision threshold offset parameter λ

(defined in [6]; $\lambda = 0$ corresponds to truncation and $\lambda = 1$ to rounding), then it can be shown that the expression for the quantizing noise power becomes

$$\frac{2}{\alpha^2} - \frac{qe^{\alpha q \left(\frac{\lambda}{2} - 1\right)}}{\alpha(1 - e^{-\alpha q})} [q\alpha(1 - \lambda) + 2]$$

This expression depends only on three quantities:

- the quantizer level spacing q , which is known from the quantizer scale code and q_scale_type parameters received in the Information Bus
- the decision threshold offset parameter λ . This is not known but it is reasonable to suppose that it takes the value 0.75 in the case of intra coding
- the Laplacian distribution parameter α . This is not known but it is possible to estimate it using one of the approaches outlined below.

We now look at the problem of estimating the parameter α . This parameter defines the sharpness of the probability distribution of the DCT coefficients. Three possible approaches have been identified:

- Direct estimation using the decoded DCT coefficients themselves. A histogram of DCT coefficients of each frequency can be built up, taking into account the quantizer scale values and the quantizer weighting matrices, both of which are known from the Information Bus. Each histogram can then be matched to a Laplacian distribution, using either a regression technique or by matching a particular parameter such as the entropy or the variance.
- Counting the number of zero DCT coefficients of each frequency. This is essentially a simplification of the first approach, and reflects the fact that the difference between the actual distributions and a uniform distribution lies chiefly in the fact that the actual distributions will have many more small values than the uniform distribution. These are values which may have quantized to zero even if the quantizer scale had been coarser.
- A simpler approach is to reduce the problem to that of estimating a single parameter across all the DCT coefficient frequencies, rather than one for each frequency. In order to do this, it is necessary to use some kind of generic relationship between the Laplacian parameters for each frequency, such as the table of typical parameters given in [6]. The single parameter then represents what would happen if coefficients with the generic distribution were scaled by a constant factor. Thus, if the generic parameters are

$$\alpha_i, i = 1 \dots 63$$

then a scaling factor k would lead to a set of parameters

$$k\alpha_i, i = 1 \dots 63$$

The single parameter k can itself be estimated in a number of ways. For example, different values of k will lead, through knowledge of the MPEG-2 variable-length code tables and of the quantizer scale and weighting matrices, to different coefficient bit rates. It follows that from the bit rate and the other information it is possible to estimate k .

It should be noted that the PSNR estimate can, by appropriate use of the quantiser weighting matrix parameters in the calculations, be made either in the weighted DCT coefficient domain or in the unweighted domain, where it would be directly linked to the pixel domain through Parseval's Theorem.

4 DCT Basis functions

One of the most annoying impairments in decoded MPEG-2 sequences is the visibility of DCT basis functions (i.e. functions that are the inverse DCT of a single non-zero coefficient). A possible method of measuring this impairment is to take the DCT of the decoded picture and record for each block the difference between the highest and second highest coefficients. The mean, or some other statistical function of this quantity such as that described in Appendix B, would provide such a measure.

5 References

1. ISO/IEC 13818-2 (MPEG-2 Video)
2. Sarnoff JND technique
3. WO 95/35628
4. WO 96/24222

Appendix A : A method for taking the picture source into account

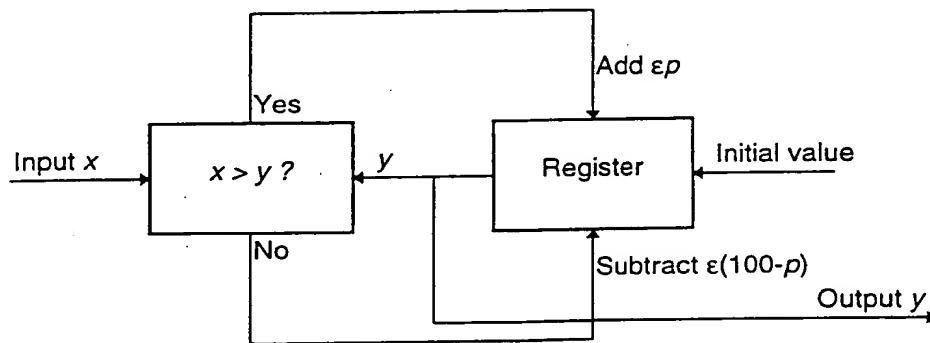
The approaches described in this paper are all based on the "single-ended" architecture of Figure 2 and as such suffer from the limitation that there is no knowledge of how much of the impairment being measured has come from the coding process and how much has come from the source. This Appendix outlines a way in which that limitation can be partially overcome.

The idea is to apply some or all of the measures to the source and/or to intermediate points in the signal chain and to transmit the results to the decoder under consideration, using a combination of ancillary data in MPEG bitstreams and the Information Bus, according to the principles of the "Video Postmark" [4]. At intermediate points in the chain, where the picture has been decoded from an MPEG bitstream and there is access to the Information Bus resulting from that decoding process, all the measures described above can be used. At the source, or at places where a full Information Bus is not available, the choice of measures may be more limited. In either case, the results can be compared with the current results and the difference will give an indication of how much of the finally measured degradation was due to the intervening compression process or processes.

Appendix B : A method for estimating centiles in probability distributions

Many methods of estimating picture quality make use of the mean of some quantity. However, as we pointed out in Section 2, it is sometimes more appropriate to use a measure such as the 90th centile of the distribution.

In order to measure accurately the p th centile of a distribution, it is necessary to build up a histogram of samples and then find the value up to which the area under the histogram is p per cent of the total area. This is a fairly complicated operation. An alternative, which gives a running estimate of the p th centile, works as shown in the following diagram:



A register is loaded with an initial estimate of the p th centile. This estimate can be obtained in any way, either through *a priori* knowledge of the problem, or simply from the first sample or few samples received. Each input sample x is then compared with the current estimate y . If it is greater, the estimate is increased by a small amount proportional to p . If it is less, the estimate is reduced by a small amount proportional to $(100 - p)$.

This means that, if the current estimate is equal to the p th centile, the ratio of comparisons yielding "greater" to those yielding "less" will be $(100 - p) : p$, so the long-term average change to the register contents will be

$p(100 - p) - (100 - p)p = 0$. If the current estimate is too large, there will be more "less" results and the register value will tend to decrease. Conversely, if the current estimate is too small, there will be more "greater" results and the register value will tend to increase. The parameter ε gives a time constant for the system.

This Page Blank (uspto)

PCHGBAA 1033 56

MATHYS + SQUIRE

11/10/99

THIS PAGE BLANK (USPTO)